Effective Combination of Language and Vision Through Model Composition and the R-CCA Method

Hagar Loeub
hagar.loeub@gmail.com

Roi Reichart
roiri@ie.technion.ac.il

Faculty of Industrial Engineering and Management, Technion, IIT

arXiv:1609.08810v2 [cs.CL] 4 Oct 2016

Abstract

We address the problem of integrating textual and visual information in vector space models for word meaning representation. We first present the *Residual CCA (R-CCA)* method, that complements the standard CCA method by representing, for each modality, the difference between the original signal and the signal projected to the shared, max correlation, space. We then show that constructing visual and textual representations and then post-processing them through *composition of common modeling motifs* such as PCA, CCA, R-CCA and linear interpolation (a.k.a *sequential modeling*) yields high quality models. On five standard semantic benchmarks our sequential models outperform recent multimodal representation learning alternatives, including ones that rely on joint representation learning. For two of these benchmarks our R-CCA method is part of the *Best* configuration our algorithm yields.

1 Introduction

In recent years, vector space models (VSMs), deriving word meaning representations from word co-occurrence patterns in text, have become prominent in lexical semantics research (Turney et al., 2010; Clark, 2012). Recent work has demonstrated that when other modalities, particularly the visual, are exploited together with text, the resulting multimodal representations outperform strong textual models on a variety of tasks (Baroni, 2016).

Models that integrate text and vision can be largely divided to two types. *Sequential models* first separately construct visual and textual representations and then merge them using a variety of techniques: concatenation (Bruni et al., 2011; Silberer et al., 2013; Kiela and Bottou, 2014), linear weighted combination of vectors (Bruni et al., 2012; Bruni et al., 2014) or linear interpolation of model scores (Bruni et al., 2014), Canonical Correlation Analysis and its kernalized version (CCA, (Hill et al., 2014; Silberer and Lapata, 2012; Silberer et al., 2013)), Singular Value Decomposition (SVD, (Bruni et al., 2014)) and Weighted Gram Matrix Combination (Reichart and Korhonen, 2013; Hill et al., 2014)). *Joint models* directly learn a joint representation from textual and visual resources using Bayesian modeling (Andrews et al., 2009; Feng and Lapata, 2010; Roller and Im Walde, 2013) and various neural network (NN) techniques: autoencoders (Silberer and Lapata, 2014), extensions of word2vec skip-gram (Hill and Korhonen, 2014; Lazaridou et al., 2015) and others (e.g. (Howell et al., 2005)).

The focused contribution of this short paper is two-fold. First, we advocate the *sequential approach* for text and vision combination and show that when a systematic search in the space of configurations of *composition of common modeling motifs* (Grosse, 2014) is employed, this approach outperforms recent *joint models* as well as sequential models that do not thoroughly search the space of configurations. This finding has important implications for future research as it advocates the development of efficient search techniques in configuration spaces of the type we explore.

Particularly, we experiment with unimodal dimensionality reduction with Principal Component

Analysis ((PCA, (Jolliffe, 2002)), multimodal fusion with Canonical Correlation Analysis (CCA, (Hardoon et al., 2004)) and model score combination with linear interpolation (LI, (Bruni et al., 2014)). The composed models outperform strong alternatives on semantic benchmarks for word pair similarity and association: MEN (Bruni et al., 2014), WordSim353 (WS, (Finkelstein et al., 2001)), SimLex999 (SL (Hill et al., 2015)), SemSim and VisSim (SSim, VSim, (Silberer and Lapata, 2014)).

Our second contribution is in proposing the *Residual CCA (R-CCA)* method for multimodal fusion. This method complements the standard CCA method by representing, for each modality, the difference between the original signal and the signal projected to the shared space. Since CCA aims to maximize the correlation between the projected signals, the residual signals intuitively represent uncorrelated components of the original signals. Empirically, including R-CCA in the configuration space improves results on two evaluation benchmarks. Moreover, for all five benchmarks R-CCA substantially outperforms CCA.

2 Multimodal Composition

2.1 Modeling Motifs

PCA is a standard dimensionality reduction method. We hence do not describe its details here and refer the interested reader to (Jolliffe, 2002).

CCA finds two projection vectors, one for each original vector, such that projecting the original vectors yields the highest possible correlation under linear projection. In short, given an n word vocabulary, with representations $X \in R^{n \times d_1}$ and $Y \in R^{n \times d_2}$, CCA seeks two sets of projection vectors $V \in R^{d_1 \times d}$ and $W \in R^{d_2 \times d}$ that maximize the correlation (ρ) between the projected vectors of each of the words: $V, W = \arg\max_{V', W'} \rho(XV', YW')$. The final projection is: $X' = XV$ and $Y' = YW$.

Residual-CCA (R-CCA) CCA aims to project the involved representations into a shared space where the correlation between them is maximized. The underlying assumption of this method is hence that multiple modalities can facilitate learning through exploitation of their shared signal. A complementary point of view would suggest that important information can also be found in the dissimilar components of the monomodal signals.

While there may be many ways to implement this idea, we explore here a simple one which we call the *residuals approach*. Denoting the original monomodal signals with X and Y and their CCA projections with X' and Y' respectively, the *residual* signals are defined as: $R_x = X - X'$ and $R_y = Y - Y'$. Notice that a monomodal signal (e.g. X) and its CCA projection (e.g. X') may not be of the same dimension. In such cases we first project the original signal (X) to the dimensionality of the projected signal (X') with PCA.

LI combines the scores produced by two VSMs for a word pair, $sc_{m1}(w_i, w_j)$ and $sc_{m2}(w_i, w_j)$, using the linear equation ($\alpha \in [0, 1]$): $Score(w_i, w_j) = \alpha \cdot sc_{m1}(w_i, w_j) + (1 - \alpha) \cdot sc_{m2}(w_i, w_j)$.

2.2 Motif Composition

We divide the above modeling motifs to three layers, to facilitate an efficient systematic optimal configuration search (Figure 1): (a) Data: (a.1) original vectors; or (a.2) original vectors projected with unimodal PCA; (b) Fusion: (b.1) CCA and (b.2) R-CCA, each method outputting two projected vectors per word, one for each modality; (c) Combination: (c.1) vector concatenation; and (c.2) linear interpolation (LI) of model scores. In our search, a higher

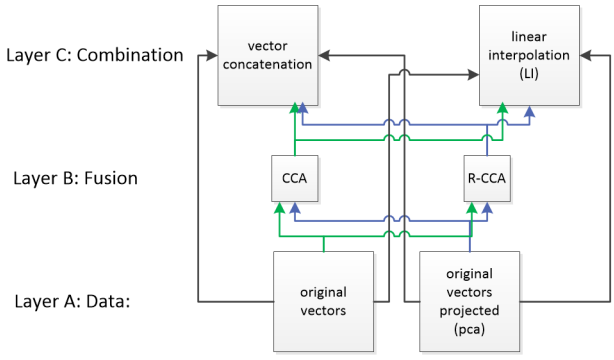

Figure 1: The motif composition search space of the sequential modeling approach for finding the best composition on pair similarity tasks.

layer method considers inputs from all lower layer methods as long as both inputs are the output of the same method. That is, CCA (layer b.1) is ap-

% pairs (in ImageNet)	MEN 42%		WS 22%		SL 29%		SSim 85%		VSim 85%	
Model	Config.	ρ	Config.	ρ	Config.	ρ	Config.	ρ	Config.	ρ
Best	PCA(200) CCA (V,200) + R-CCA (T,200) LI (0.4)	**0.81**	PCA (250) CCA (250) LI (0.95)	**0.74**	PCA (250) CCA (V,250) + R-CCA (T,250) LI (0.6)	**0.62**	PCA (50) CCA (no) LI (0.7)	**0.8**	PCA (100) CCA (no) LI (0.45)	**0.68**
Best (No R-CCA)	PCA(100) CCA (50) LI (0.65)	0.79			PCA(150) CCA (150) LI (0.4)	0.61				
Best (LI)	0.7	0.79	0.95	0.71	0.45	0.56	0.8	0.75	0.55	0.66
Best (PCA, Skip)	350	0.68	250	0.55	300	0.54	100	0.73	100	0.67
Best (PCA, CNN)	50	0.63	100	0.52	300	0.53	50	0.7	100	0.66
Best (CCA)	V, 100	0.59	T, 150	0.61	V, 50	0.47	T, 50	0.39	V, 100	0.35
Best (R-CCA)	L, 50	0.77	L, 150	0.71	V, 300	0.53	L, 50	0.75	V, 50	0.66
Concatenation	–	0.63	–	0.48	–	0.54	–	0.57	–	0.6
MMSKIP-A	–	0.74	–	–	—	0.5	–	0.72	–	0.63
MMSKIP-B	–	0.76	–	–	—	0.53	–	0.68	–	0.6
BR-EA-14	–	0.77	–	–	—	0.44	–	0.69	–	0.56
KB-14	–	0.74	–	0.57	—	0.33	–	0.60	–	0.50
Skip	–	0.75	–	0.73	–	0.46	–	0.73	–	0.67
CNN	–	0.58	–	0.54	–	0.53	–	0.67	–	0.64

Table 1: Results. *B*est is the best configuration for each benchmark. *Best (*)* and *Concatenation* are the best single-motif models. MMSKIP-A and MMSKIP-B are the (joint) models of Lazaridou et al. (2015), BR-EA-14 is the best performing model of Bruni et al. (2014) and KB-14 is the model of Kiela and Bottou (2014) (both are sequential). If LI is employed after CCA then its input are the CCA output vectors. If it is employed after a CCA+R-CCA then one of its input vectors comes from CCA and the other from R-CCA. For PCA and CCA we report the number of dimensions and for LI the weight of the textual model. For CCA we also report whether the textual (T) or the visual (V) projected vector yielded the best result. *Skip* and *CNN* are the original input vectors.

plied to original textual and visual vector pairs (output of a.1) as well as to PCA-transformed vectors (a.2), but not, e.g., to PCA-transformed visual vectors (a.2) paired with original textual vectors (a.1). Vector concatenation (c.1) and linear score interpolation (c.2), in turn, are applied to all the inputs and outputs of CCA and of R-CCA.

The only exception is that we allow an output of CCA (e.g. projected visual vectors) and an output of R-CCA (e.g. residual textual vectors) as input to layer c, as two projections or two residuals may convey very similar information. To facilitate efficiency further, CCA and R-CCA are only applied to textual and visual vectors of the same dimensionality. We leave the exploration of other, possibly more complex, search spaces to future work.

For each benchmark (Section 3) we search for its **Best** configuration: the optimal sequence of the above motifs, **at most** one from each layer, together with the optimal assignment of their parameters. We do not aim to develop efficient algorithms for optimal configuration inference, but rather employ an exhaustive grid search approach. The high quality configurations we find, advocate future development of efficient search algorithms.

3 Data and Experiments

Input Vectors Our textual VSM is word2vec skip-gram (Mikolov et al., 2013), [1] trained on the 8G words corpus generated by the word2vec script.[2] We followed the hyperparameter setting of (Schwartz et al., 2015) and, particularly, set vector dimensionality to 500. For the visual modality, we used the 5100 4096-dimensional vectors of Lazaridou et al. (2015), extracted with a pre-trained Convolutional Neural Network (CNN, (Krizhevsky et al., 2012)) and the Caffe toolkit (Jia et al., 2014) from 100 pictures sampled for each word from its ImageNet (Deng et al., 2009) entry. While there are various alternatives for both textual and visual representations, those we chose are based on state-of-the-art techniques.

Benchmarks We report the Spearman rank correlation (ρ) between model and human scores, for the word pairs in five benchmarks: MEN, WS, SL, SSim and VSim. While all the words in our benchmarks appear in our textual corpus, only a fraction of them appears in ImageNet, our source of visual input. Hence, following Lazaridou et al. (2015),

[1] https://code.google.com/p/word2vec/

[2] code.google.com/p/word2vec/source/browse/trunk/demo-train-big-model-v1.sh

for each benchmark we report results only for word pairs consisting of words that are represented in ImageNet. A model word pair score is the cosine similarity between the vectors learned for its words.

Parameter Tuning We jointly optimized parameters together with the decision of which modeling motif to select at each layer, if at all. For PCA, CCA and R-CCA we iterated over dimensionality values from 50 onward in steps of 50, till the minimum dimensionality of the input sets. For LI, we iterated over $\alpha \in \{0, 0.1, \ldots, 1\}$. Among the best performing configurations we selected the one with the lowest dimension output vectors.

Note that there is no agreed upon split of our benchmarks, except from MEN, to development and test portions. Therefore, to facilitate future comparison with our work, the main results we report are with the best configuration for each benchmark, as tuned on the entire benchmark. We also show that our results generalize well across evaluation sets, including MEN's dev/test split.

Alternative Models We compare our results to strong alternatives: MSKIP-A and MMSKIP-B ((Lazaridou et al., 2015), joint models), the best performing model of Bruni et al. (2014) and Kiela and Bottou (2014) (sequential models).[3] While our results are not directly comparable to these models due to different training sources and parameter tuning strategies,[4] this comparison puts them in context.

4 Results

Table 1 presents the results. *Best* outperforms the unimodal models (*S*kip and *C*NN) and the alternative models. The gains (in ρ points) are: MEN: 4, WS: 1, SL: 9, SSim: 7 and VSim: 1. R-CCA is included in *Best* for MEN and SL, improving over the best configuration that does not include it by 2 and 1 ρ points, respectively. Furthermore, R-CCA outperforms CCA on all five benchmarks (by 6-31 ρ points) and particularly on MEN, SSim and VSim.

Observations The five *Best* configurations share meaningful patterns. **(1)** *Best* never includes concatenation (layer c.1); **(2)** LI is always included in *Best* and the weights assigned to the textual and visual modalities are mostly balanced. Particularly, the weight of the textual modality is 0.4-0.7 for MEN, SL, SSim and VSim; **(3)** In 3 out of 5 cases, *B*est (LI), that linearly interpolates the scores yielded by the input textual and visual vectors without PCA, CCA or R-CCA processing, outperforms the models from previous work and the unimodal models; **(4)** In all *Best* configurations, the reduced dimensionality is 50-250, which is encouraging as processing smaller vectors requires less resources.

Generalization We now show that our results generalize well across evaluation sets. First, for the portion of the MEN development set that overlaps with ImageNet, our *Best* MEN configuration is the third-best configuration, with $\rho = 0.78$. We also tested the *Best* configuration as tuned on each of the benchmarks, on the remaining four benchmarks. We observed that WS, MEN and SL serve as good development sets for each other. Best SL configuration (*Best-SL*): $\rho = 0.7$ on WS and $\rho = 0.77$ on MEN, *Best-WS*: $\rho = 0.61$ (SL) and $\rho = 0.74$ (MEN), and *Best-MEN*: $\rho = 0.6$ (SL) and $\rho = 0.71$ (WS). The performance of these models on SSim and VSim, however, is substantially lower than that of the *Best* models of these sets (e.g. $\rho = 0.62$ for *Best-SL* on SSim, $\rho = 0.5$ for *Best-WS* on VSim, compared to of $\rho = 0.8$ and $\rho = 0.68$, respectively). Likewise, SSim and VSim, that consist of the same word pairs scored along different dimensions, form good development sets for each other ($\rho = 0.66$ for *Best-SSim* on VSim, $\rho = 0.78$ for *Best-VSim* on SSim), but not for WS or SL. That is, each benchmark has other benchmarks that can serve as its dev. set.

5 Conclusions

We demonstrated the power of *composition of common modeling motifs* in multimodal VSM construction and presented the R-CCA method that exploits the residuals of the CCA signals. Our model yields

[3] To facilitate clean comparison with previous work, we copy the results of Bruni et al. (2014) and of Kiela and Bottou (2014) from Lazaridou et al. (2015), except for WS. Lazaridou et al. (2015) do not report results for WS, while Bruni et al. (2014) report results on a different subset than ours, consisting of 252 word pairs. Kiela and Bottou (2014) report results on our subset of WS, and we copy their best result. Silberer and Lapata (2014) also report results for SSim and VSim but for the entire sets rather than for our subsets.

[4] Section 5 of Lazaridou et al. (2015) provides the details of the alternative models, their training and parameter tuning.

state-of-the-art results on 5 leading semantic benchmarks, for two of which R-CCA is part of the *Best* configuration. Moreover, R-CCA performs much better than CCA on all five benchmarks.

Our results hence advocate two research directions. First, they encourage *sequential modeling* with systematic search in the configuration space for multimodal combination. Our future goal is making model composition a standard tool for this problem, by developing *efficient inference algorithms* for optimal configurations in possibly more complex search spaces than those we explored with an exhaustive grid search. Second, the encouraging results of R-CCA emphasize the potential of informed post-processing of the CCA output. We intend to deeply delve into this issue in the immediate future.

References

[Andrews et al.2009] Mark Andrews, Gabriella Vigliocco, and David Vinson. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological review*, 116(3):463.

[Baroni2016] Marco Baroni. 2016. Grounding distributional semantics in the visual world. *Language and Linguistics Compass*, 10(1):3–13.

[Bruni et al.2011] Elia Bruni, Giang Binh Tran, and Marco Baroni. 2011. Distributional semantics from text and images. In *Proc. of the GEMS workshop on geometrical models of natural language semantics,EMNLP*, pages 22–32.

[Bruni et al.2012] Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proc. of ACL*, pages 136–145.

[Bruni et al.2014] Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research (JAIR)*, 49:1–47.

[Clark2012] Stephen Clark. 2012. Vector space models of lexical meaning. *Handbook of Contemporary Semantics, Wiley-Blackwell, à paraître*.

[Deng et al.2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Proc. of CVPR*, pages 248–255.

[Feng and Lapata2010] Yansong Feng and Mirella Lapata. 2010. Visual information in semantic representation. In *Proc. of NAACL*, pages 91–99.

[Finkelstein et al.2001] Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. In *Proc. of WWW*, pages 406–414.

[Grosse2014] Roger Baker Grosse. 2014. *Model selection in compositional spaces*. Ph.D. thesis, Massachusetts Institute of Technology.

[Hardoon et al.2004] David Hardoon, Sandor Szedmak, and John Shawe-Taylor. 2004. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, 16(12):2639?–2664.

[Hill and Korhonen2014] Felix Hill and Anna Korhonen. 2014. Learning abstract concept embeddings from multi-modal data: Since you probably can't see what i mean. In *Proc. of EMNLP*, pages 255–265.

[Hill et al.2014] Felix Hill, Roi Reichart, and Anna Korhonen. 2014. Multi-modal models for concrete and abstract concept meaning. *Transactions of the Association for Computational Linguistics*, 2:285–296.

[Hill et al.2015] Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.

[Howell et al.2005] Steve R Howell, Damian Jankowicz, and Suzanna Becker. 2005. A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language*, 53(2):258–276.

[Jia et al.2014] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proc. of the ACM International Conference on Multimedia*, pages 675–678. ACM.

[Jolliffe2002] Ian Jolliffe. 2002. *Principal component analysis*. Wiley Online Library.

[Kiela and Bottou2014] Douwe Kiela and Léon Bottou. 2014. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proc. of EMNLP*, pages 36–45.

[Krizhevsky et al.2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Proc. of NIPS*.

[Lazaridou et al.2015] Angeliki Lazaridou, Nghia The Pham, and Marco Baroni. 2015. Combining language and vision with a multimodal skip-gram model. In *Proc. of NAACL*.

[Mikolov et al.2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proc. of NIPS*.

[Reichart and Korhonen2013] Roi Reichart and Anna Korhonen. 2013. Improved lexical acquisition through

dpp-based verb clustering. In *Proc. of ACL*, pages 862–872.

[Roller and Im Walde2013] Stephen Roller and Sabine Schulte Im Walde. 2013. A multimodal lda model integrating textual, cognitive and visual modalities. In *Proc. of EMNLP*, pages 1146–1157.

[Schwartz et al.2015] Roy Schwartz, Roi Reichart, and Ari Rappoport. 2015. Symmetric pattern based word embeddings for improved word similarity prediction. In *Proc. CoNLL*.

[Silberer and Lapata2012] Carina Silberer and Mirella Lapata. 2012. Grounded models of semantic representation. In *Proc. of EMNLP-CoNLL*.

[Silberer and Lapata2014] Carina Silberer and Mirella Lapata. 2014. Learning grounded meaning representations with autoencoders. In *Proc. of ACL*.

[Silberer et al.2013] Carina Silberer, Vittorio Ferrari, and Mirella Lapata. 2013. Models of semantic representation with visual attributes. In *Proc. of ACL*, pages 572–582.

[Turney et al.2010] Peter D Turney, Patrick Pantel, et al. 2010. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37(1):141–188.